Published in *Scientific Reports* **1**, 151 (2011); DOI:10.1038/srep00151

arXiv:1111.6115v1 [physics.soc-ph] 25 Nov 2011

Discovering Network Structure Beyond Communities

Takashi Nishikawa[1,*] and Adilson E. Motter[2,3]

[1]*Department of Mathematics, Clarkson University, Potsdam, NY 13699, USA*
[2]*Department of Physics and Astronomy and Northwestern Institute on Complex Systems, Northwestern University, Evanston, IL 60208, USA*
[3]*Department of Molecular Biology, Princeton University, Princeton, NJ 08544, USA*

To understand the formation, evolution, and function of complex systems, it is crucial to understand the internal organization of their interaction networks. Partly due to the impossibility of visualizing large complex networks, resolving network structure remains a challenging problem. Here we overcome this difficulty by combining the visual pattern recognition ability of humans with the high processing speed of computers to develop an exploratory method for discovering groups of nodes characterized by common network properties, including but not limited to communities of densely connected nodes. Without any prior information about the nature of the groups, the method simultaneously identifies the number of groups, the group assignment, and the properties that define these groups. The results of applying our method to real networks suggest the possibility that most group structures lurk undiscovered in the fast-growing inventory of social, biological, and technological networks of scientific interest.

The highly structured internal organization of complex networks can both impact and reflect their dynamics and function[1]. Previous work on identifying and studying this organization has focused mainly on network communities[2–13], which are subsets of nodes defined by the difference between their internal and external link density. To provide a fresh perspective on this problem, we seek to capture more general structures characterized by other network properties[14–17]. For this purpose, we introduce the notion of *structural groups*, defined as subsets of nodes sharing common structural properties that set them apart from other nodes in the network. Using a given set of $p \gg 1$ node properties (such as centrality and spectral properties) as the coordinates for each node in the p-dimensional space $\mathbb{R}^p$, we identify structural groups as clusters of points in this *node property space*. Figure 1 shows an illustrative example of a network for which no standard network visualization shows clear group structure (Fig. 1**a**). However, an appropriate two-dimensional projection in the node property space reveals a hidden, but unambiguous three-group structure (Fig. 1**b**), which can be used to generate a far more informative layout of the network (Fig. 1**e**). Application of existing community detection methods[18–27] is not expected to resolve these groups, since they are not distinguishable by link density alone (Fig. 1**c**). Neither is the direct application of existing clustering methods in the full node property space nor in the projection onto any lower-dimensional space, due to the known fact that groups with widely different scatter sizes may not be correctly grouped by unsupervised algorithms (Fig. 1**d** and Supplementary Fig. S1). Distinguishing structural groups may in general require a combination of two or more properties — Fig. 1**b** shows that the degree and the average degree of neighbors suffice for this example. It is difficult, however, to identify such a combination without knowing the groups a priori.

Our approach overcomes these difficulties using the visual processing ability of a human user as an integral part of the analysis. The approach is based on visual analytics[28,29], which is conceptualized as exploratory statistics in which analytical reasoning is facilitated by a visual interactive interface. Humans generally excel automated computer algorithms in visual recognition tasks, such as labeling images[30] and deciphering distorted texts, which forms the basis of spam prevention systems and crowdsourcing for the digitalization of old books[31]. We exploit this capability by asking the user to inspect a selection of two-dimensional projections of the node property space for possible separation of nodes into groups. Since any projection could potentially reveal good separation of groups, we first consider the result of choosing these projections randomly. For two clusters of points with a gap between them in high dimension, the probability can be very small for the clusters to be separable by a straight line in a random two-dimensional projection. This probability depends strongly on the "effective dimension" of the clusters. For example, if two Gaussian-distributed clusters of 100 points have their centers 6 units apart in the 28-dimensional space, the probability is less than 0.001 if the variance of the clusters in every direction is one, but increases to about 0.017 if the variance is reduced by a factor of 10 in all but 10 orthogonal directions. We find that the effective dimension is relatively small for the groups discovered in the networks considered here, most of them with dimension less than 12 (out of 28) when defined as the minimum number of principal components required to account for 90% of the variance within the group. To further enhance the probability of separating groups, we sample random projections with a systematic bias (see Methods). This increases the sep-

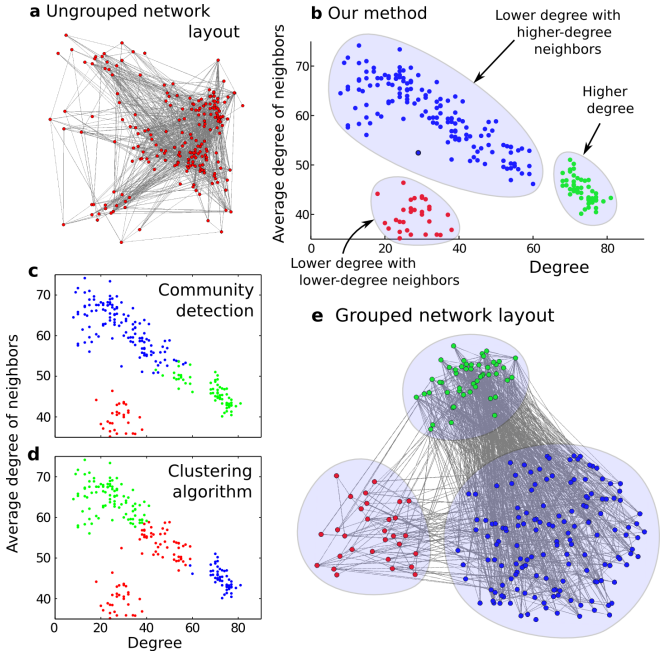

Figure 1: **Discovering hidden group structure beyond density-based communities. a**, Visualization of a network by the Gürsoy-Atun algorithm[51], which attempts to place nodes uniformly while keeping the network neighbors close. This and other standard layout algorithms fail to disentangle the network and reveal any clear group structure. **b**, Using our visual analytics method, a user can discover three structural groups (of sizes 150, 50, and 30) without a priori information about the number of groups. The groups can be characterized by the degree and the neighbors' average degree, and at least two properties are necessary to resolve the entire group structure. **c**, Even the most general community detection method[14] does not divide the network correctly. **d**, The K-means algorithm[45], one of the most frequently used methods for general clustering problems, does not correctly capture the group structure when applied directly to the full node property space, even if the number of groups $K = 3$ is given. **e**, Layout of the network using the discovered groups. For clarity, both panels **a** and **e** show only 10% of the links.

aration probability for the example of Gaussian clusters above to around 0.68 for a single projection. If the user visually recognizes separation of nodes into groups in a two-dimensional projection, the group assignment is entered through a graphical interactive interface (Fig. 2**a**–**d**, Supplementary Video S1). The integration of the visual component allows the user not only to supervise the process, but also to learn and create intuition from taking part in the process, thus facilitating the search for unanticipated network structures. It also accommodates naturally an ultimate goal of clustering algorithms, which is to reproduce how a human would group a given set of points.

The chance of capturing a group structure is even further enhanced by the multiplicative effect of using more than one projection. Indeed, the separation probability in the example above rises from 0.68 to above 0.999 with just 7 projections. In general, for a given number L of random projections, the probability that all of these projections fail to separate a given pair of group decreases to zero exponentially with L. After the user processes a given number L of projections, each node i in the network will be associated with a group assignment vector $\mathbf{a}^{(i)}$ representing the user input (Fig. 2**d**). Since we typically have a large number of distinct assignment vectors, we aggregate the corresponding nodes into a smaller, more meaningful number of structural groups by single-linkage hierarchical clustering[32]. For this, we use the Hamming distance between the group assignment vectors of different nodes, $\mathbf{a}^{(i)}$ and $\mathbf{a}^{(j)}$, which in this case is the number of projections for which the user has placed those nodes in different groups. This results in a dendrogram that we can cut at a threshold distance d to obtain a grouping, in which being in different groups indicates that the user has placed these nodes in different groups in at least d out of L projections (Fig. 2**e**; Supplementary Video S1). To compare the different groupings obtained at different thresholds, we define the *quality of grouping* Q_g by

$$Q_g = \frac{1}{q_g} \cdot \frac{\langle ||\mathbf{c}_k - \mathbf{c}_\ell|| \rangle_{k,\ell}}{\langle ||\mathbf{x}^{(i)} - \mathbf{c}_{k_i}|| \rangle_i}, \qquad (1)$$

where vector $\mathbf{x}^{(i)} = (x_1^{(i)}, x_2^{(i)}, \ldots, x_p^{(i)})^T$ represents node i in the property space, vector $\mathbf{c}_k$ is the center of group k, index k_i denotes the group to which node i belongs, and $||\cdot||$ defines the p-dimensional Euclidean distance. The ratio of the two bracketed quantities in Eq. (1) measures the average separation distance between groups (the average over all pairs of groups, denoted $\langle\cdot\rangle_{k,\ell}$) relative to the spread within individual groups (the average over all nodes, denoted $\langle\cdot\rangle_i$). This quantity is then normalized by a constant q_g, chosen to remove a systematic dependence of the quality of grouping on the number of groups K (see Methods). As one lowers the threshold level, the quality of grouping Q_g tends to drop sharply at a certain level (Fig. 2**f**). To obtain the maximum number of high-quality groups, we suggest choosing the group assignment, as well as the number of groups K, at the threshold level just above the largest drop in Q_g, which we call the Q_g *drop-off*.

Results

We implemented our *visual analytics method* using the selection of $p = 28$ node properties listed in Table I, which encompasses important node attributes that capture local information, such as degree and clustering, and others that capture more global information, such as betweenness centrality and Laplacian eigenvectors. In particular, the eigenvectors of the Laplacian and of the normalized Laplacian allow the detection of communities[12,33–36] and bipartite or multipartite structures[37], respectively, as well as mixtures of these structures, assuring our method the ability to detect group structures defined by link density as special cases. Using this set of properties for the

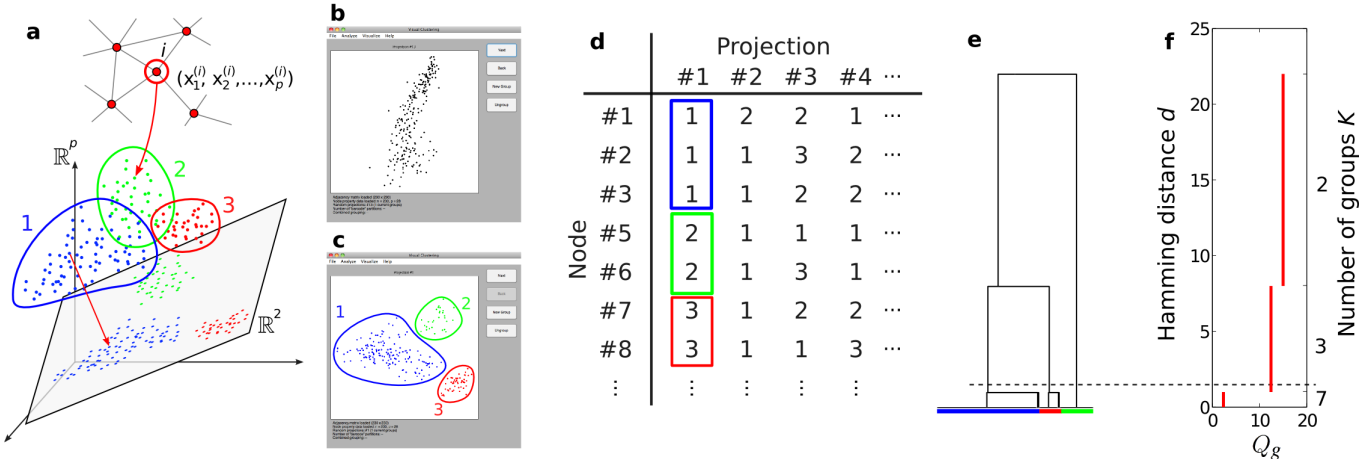

Figure 2: **Our visual analytics method.** **a**, The p node properties $x_1^{(i)}, \ldots, x_p^{(i)}$ are computed for each node i in a given network of n nodes. The nodes are then represented as points in $\mathbb{R}^p$, which are projected onto a randomly chosen two-dimensional subspace. **b**, **c**, Using a graphical interface, the user can either reject the projection (**b**), which indicates that there is no visible group separation, or indicate visible groups (**c**), which automatically assigns a group index to each node for that particular projection. **d**, Repeating this for a given number of random projections, each node i is associated with a group assignment vector $\mathbf{a}^{(i)}$, listing the group indices the user has assigned to node i. We used 30 projections for all results in this article. **e**, Dendrogram obtained by clustering the vectors $\mathbf{a}^{(1)}, \ldots, \mathbf{a}^{(n)}$. Cutting the dendrogram at a threshold Hamming distance d produces a grouping for the network. **f**, Quality of grouping Q_g as a function of the threshold level d. The appropriate number of groups is determined to be $K = 3$ by thresholding at the Q_g drop-off (dashed line).

example network of Fig. 1, we obtain the dendrogram shown in Fig. 2**e**. The number of groups for this network is found to be $K = 3$ at the Q_g drop-off (Fig. 2**f**), which agrees with the group separation visible in the projection shown in Fig. 1**b**. This accurately reflects the fact that the network was synthetically constructed from three distinct structural groups: the first two groups characterized by high (≥ 65) and low (≤ 55) prescribed degrees, respectively, but connected randomly otherwise, and the third group characterized by higher connection probability with internal nodes (0.3) than with external ones (0.1). This example illustrates that our method is capable of discovering not only group structures defined by link density, but also more general group structures, even when different types of structures coexist in the same network. Moreover, as shown in Fig. 3 for two-group benchmark networks, the visual analytics method is generally expected to outperform existing methods if the groups have different internal structures, in this case determined by their different degree distributions (see Methods).

Figure 4 shows a visualization of the hierarchy of nested structural groups identified by applying our method to a selection of six real-world networks spanning different sizes and domains (Table II). To further characterize these groups, we rank the node properties based on a two-dimensional projection in which the discovered groups reveal maximal separation (see Methods). We then discard the low-ranking properties that have negligible effect on the group separation, keeping only those indicated under each panel. Surprisingly, while most groups cannot be identified using a single node property, the node structural groups are completely separated in

Table I: **Node properties used to generate our results.**

j	$x_j^{(i)}$ (jth property of node i)
1, 2	The degree[a] of node i and the average degree of the neighbors of node i
3	The clustering coefficient[b] of node i
4	The average shortest path length from node i to all the other nodes
5, 6	The betweenness centrality[c] of node i and the average of the same quantity over the neighbors of node i
7, 8	The subgraph centrality[d] of node i and the average of the same quantity over the neighbors of node i
9–18	The ith component of the eigenvector associated with the 2nd (smallest nonzero) through the 11th eigenvalue of the Laplacian matrix[e]
19–28	The ith components associated with the 10 largest eigenvalues of the normalized Laplacian matrix[f]

Here we consider only undirected and unweighted networks for simplicity. Each quantity is normalized to the unit interval $[0, 1]$ before applying our analysis, which in this case reduces the node property space to the 28-dimensional unit hypercube.
[a] The number of links attached to node i.
[b] The fraction of pairs of neighbors of node i that are connected.
[c] The number of shortest paths passing through node i.
[d] The weighted sum of the number of closed paths in which node i participates[39].
[e] The Laplacian matrix L is defined by $L_{ij} = -1$ if nodes i and $j \neq i$ are connected, $L_{ij} = 0$ if they are not connected, and L_{ij} equals the degree of node i if $i = j$.
[f] The normalized Laplacian matrix is obtained by dividing each L_{ij} by the degree of node i.

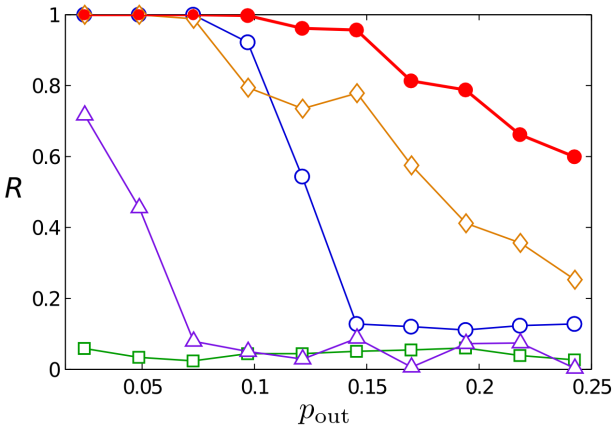

Figure 3: **Performance comparison for detecting density-based communities.** Using benchmark networks consisting of two groups, we compare the performance of the visual analytics method against alternative methods, measured by the adjusted Rand index[38] R between the computed and the true groupings (see Methods for the details of our benchmarking procedure). Our method (red filled circles) finds the correct group assignment almost perfectly for inter-group connection probability $p_{\text{out}} \lesssim 0.15$, and performs reasonably well for larger values of p_{out}. The mixture model method[14] (blue open circles) performs well for $p_{\text{out}} \lesssim 0.10$. The K-means algorithm[45] (green open squares) shows consistently low performance. The use of nonlinear kernels, dimensionality reduction based on principal component analysis, and alternative schemes for assigning weights for node properties (see Methods) led to improved performance only for small p_{out} values. One of the largest such improvement is shown here (purple open triangles). In contrast, replacing the human user with the K-means algorithm to process the two-dimensional projections in our method (see Methods) shows significantly better performance (orange open diamonds) than the direct application of the K-means variants to the node property space (green open squares and purple open triangles), although still worse than the visual analytics method (red filled circles). This demonstrates both the effectiveness of the multiple random projection approach and the advantage of the human interactive component over unsupervised algorithms. Each point in the plot is the average of R computed after removing two outliers (smallest and largest R) from a total of 20 network realizations. The visual analytics method is generally expected to outperform existing methods if the groups have different internal structures, in this case determined by their different degree distributions (see Methods).

this plane for four of the networks. The groups in three of the networks, the polbooks, netscience, and disease networks (Fig. 4**d-f**), are separated using two eigenvectors of the Laplacian matrix, suggesting that these groups could be similar to density-based communities detected by existing methods[2]; when quantified by the Rand index[38], however, the similarity appears relatively low (Supplementary Fig. S2). The groups in a fourth network, the karate network (Fig. 4**a**), can also be separated in a plane, but this projection requires the use of 15 properties led

Table II: **Networks analyzed.**

Dataset	n	m	K	d_H	p_d
karate	34	78	3	3	15
polbooks	105	441	2	2	2
adjnoun	112	425	4	2	5
football	115	613	7	3	3
netscience	379	914	4	4	2
disease	516	1188	2	5	2

We used the following datasets: **karate**, the social network of Zachary's Karate club[40]; **polbooks**, a network of political books[49]; **adjnoun**, a network of English words[36]; **football**, a network of collegiate American football teams[2]; **netscience**, a network of network scientists[36]; and **disease**, a human disease network[50] (see Supplementary Table SI for a more detailed description of these networks). Here n is the number of nodes, m is the number of links, K is the number of structural groups at the Q_g drop-off, d_H is the Hamming distance at the Q_g drop-off, and p_d is the minimum number of node properties necessary to produce a two-dimensional projection in which most or all of the K discovered groups are visually separated.

by the average degree, average betweenness, and average subgraph centrality[39] of neighbors (see Table I notes for the definition). The groups in the other two networks, the adjnoun and football networks (Fig. 4**b-c**), are mostly but not completely separated in this two-dimensional representation. We emphasize that it is not necessary for all the groups to be separable in a single two-dimensional projection. In fact, while each such projection may only illuminate part of the hidden group structure (such as the separation between a single group and all the others), the multiplicative effect of integrating information from many random projections is what often reveals the full high-dimensional structure.

Another remarkable feature of this approach is that, because we do not know in advance which properties define the groups we seek to identify, the visual analytics method simultaneously provides the answer to the question—the number and identity of the structural groups—along with the question itself—the properties that define these groups. Even when these properties are abstract, further analysis can easily reveal the nature of the network's internal organization. For example, consider the karate network, whose nodes are members of a karate club and links are interactions between two members in at least one context external to the club activities. The three structural groups identified in Fig. 4**a** correspond to (1) members who are central to the club and interact with many other members; (2) peripheral members interacting only with very few, but central members; and (3) members forming a community connected to the rest of the network only through one central member (Supplementary Fig. S3). Incidentally, one of the groups we identify consists of nodes that are connected to those outside the group but to none within the group. This social group structure is markedly different from the well-studied eventual split of the club into two clubs[40].

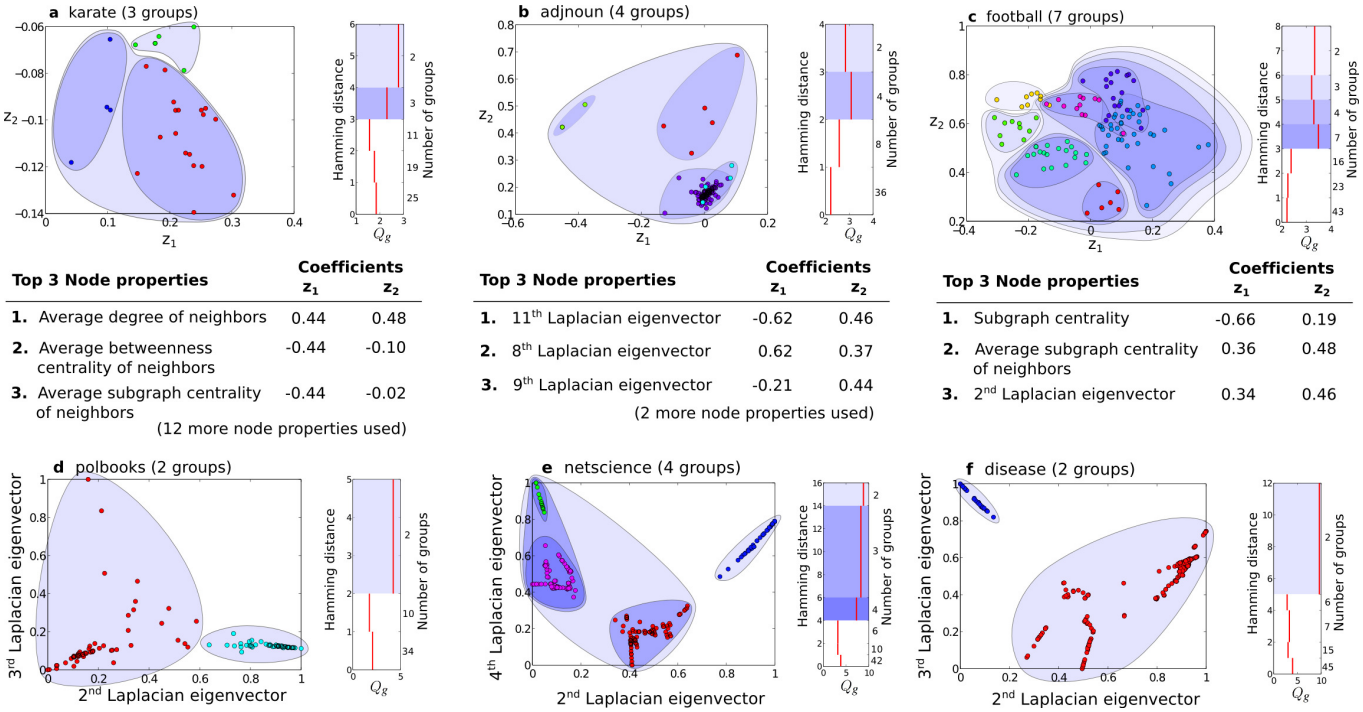

Top 3 Node properties	Coefficients z_1	z_2
1. Average degree of neighbors	0.44	0.48
2. Average betweenness centrality of neighbors	-0.44	-0.10
3. Average subgraph centrality of neighbors	-0.44	-0.02

(12 more node properties used)

Top 3 Node properties	Coefficients z_1	z_2
1. 11th Laplacian eigenvector	-0.62	0.46
2. 8th Laplacian eigenvector	0.62	0.37
3. 9th Laplacian eigenvector	-0.21	0.44

(2 more node properties used)

Top 3 Node properties	Coefficients z_1	z_2
1. Subgraph centrality	-0.66	0.19
2. Average subgraph centrality of neighbors	0.36	0.48
3. 2nd Laplacian eigenvector	0.34	0.46

Figure 4: **Hierarchical group structures discovered for the networks in Table II.** Each panel shows the network nodes plotted in a two-dimensional projection of the node property space. For panels **a**–**c**, the two coordinate axes, z_1 and z_2, are linear combinations of a selection of node properties that capture most of the group separation, with corresponding coefficients listed in the table below each plot. Note that in panels **b** and **c** two dimensions are not sufficient to cleanly separate the groups, even though our method resolves this separation by combining multiple two-dimensional projections. For panels **d**–**f**, only two properties are necessary to resolve the groups clearly. The plot on the right of each panel shows the quality of grouping Q_g as a function of the hierarchical level measured by the Hamming distance. The groupings corresponding to the hierarchical levels in the blue part of this plot are indicated in the projections by the shades of blue.

As an additional example, consider the football network, where nodes are college American football teams and links indicate matches played in the 2000 season. Although the teams are organized into 12 conferences (including Independents), our method identifies 7 structural groups (Fig. 4**c**). As shown in Figs. 5**a** and 5**b**, groups 1 and 6 are characterized by the combination of high degrees, high subgraph centrality, and the same characteristics for their neighbors, while these two groups are distinct in clustering coefficient and some Laplacian eigenvectors. Low degrees and low subgraph centrality, as well as the same characteristics for the neighbors, distinguish groups 4 and 7 from others, while they differ in their clustering coefficient and a few Laplacian eigenvectors. Group 2 shows similar characteristics as group 1 in terms of subgraph centrality, but the mean shortest path distance is very high and the betweenness centrality of the neighbors is very low, reflecting the peripheral location of these nodes within the network. Many of the Laplacian eigenvectors contribute to the separation of the groups, which is consistent with the fact that a density-based community structure exists in addition to other group structures. In particular, groups 3 and 5 are communities that can only be distinguished by the differences in the Laplacian eigenvectors and clustering coefficient. Grouping together Big Twelve and Mountain West as well as Atlantic Coast and Big East, but splitting the Independents (Fig. 5**c**), this group structure captures a higher-level organization of the conferences which is determined by the geographic proximity of the teams (Fig. 5**d**). Similar geographical manifestation of network communities has recently been observed in the effective boundaries defined by human mobility in the US[41] and telecommunications in Great Britain[42].

Discussion

The structural groups identified by the visual analytics method are characterized by common network properties. This provides a foundation for the study of the interplay between form and function in complex networks, as network dynamics (and hence function) is believed to be strongly influenced by network structure. The possibilities are extensive with our approach since the user has complete freedom to choose the set of p node properties. Within the wide range of possible structures expressible through these properties, the visual analytics method can help discover a specific group structure of interest and interpret it using a ranking of the node properties. The

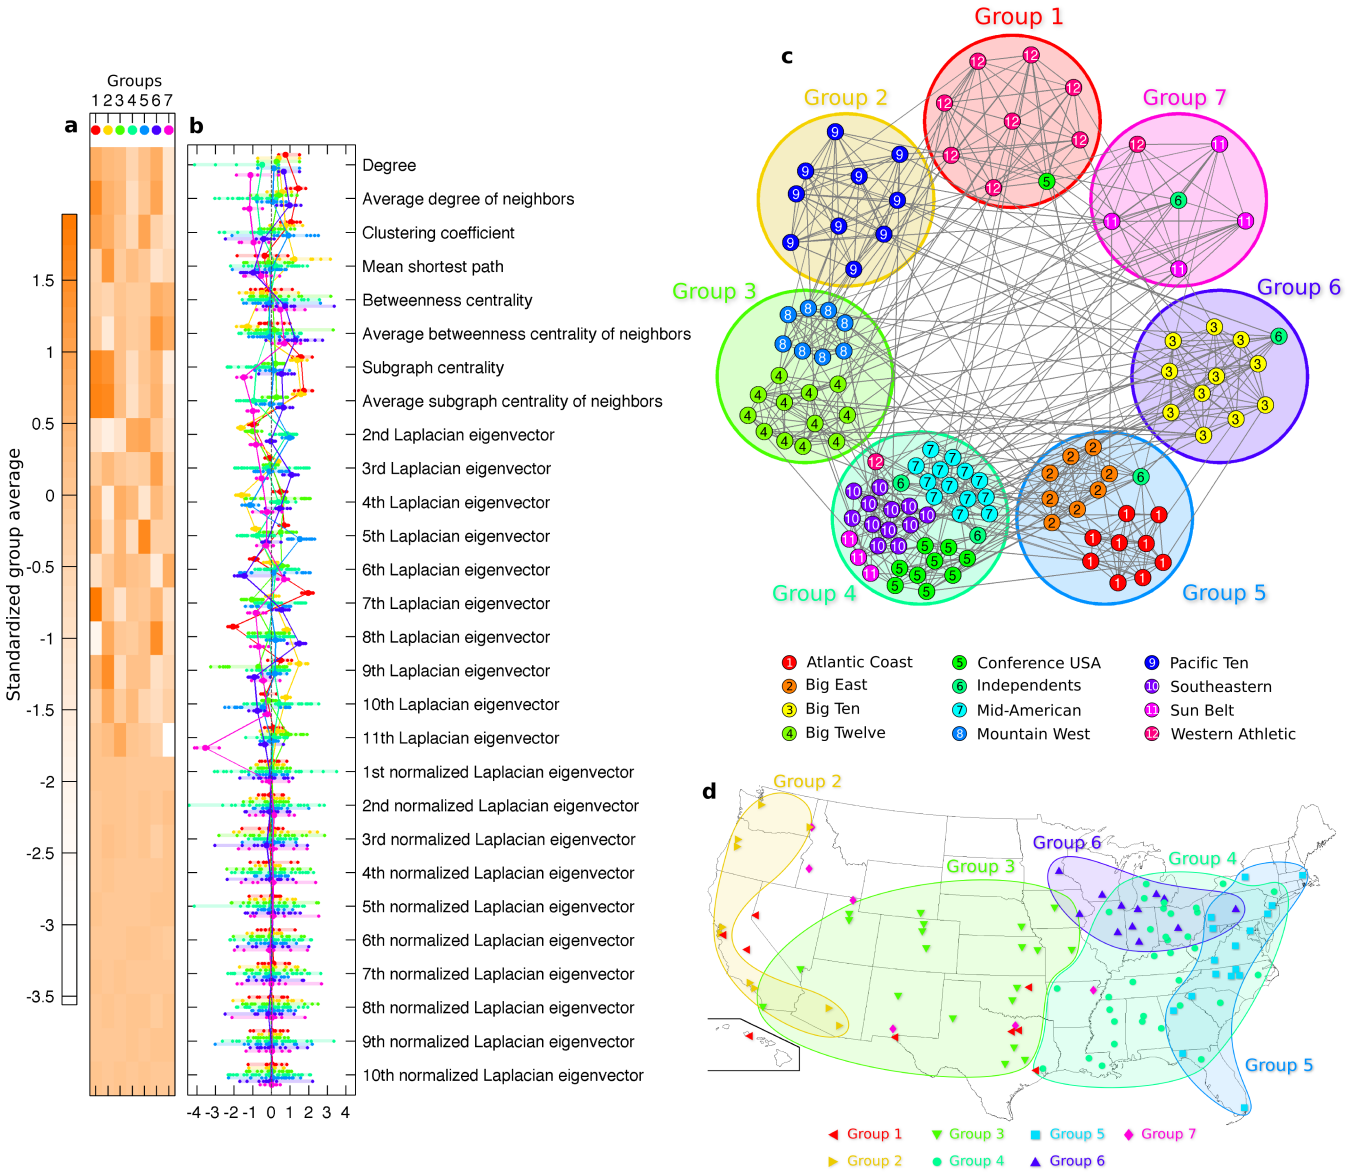

Figure 5: **Characterizing seven structural groups discovered in the football network. a**, Average node properties of the seven groups. Rows correspond to node properties and columns to groups (the colored disks at the top). Using the orange color-scale on the left, each cell shows the average node property of the group, relative to the network average and in units of the network standard deviation. **b**, Node property distribution within each group. The seven groups in this plot are color-coded as in the disks at the top of panel **a**. Small dots indicate the individual values for each node in the network, larger dots connected by lines indicate the group averages, and bars indicate the range of values for each group. All values are measured relative to the network average and in units of the network standard deviation. **c**, Layout of the network with the structural groups indicated by circles, color-coded as in the other panels. The number and color on a node indicate the college football conference to which the corresponding team belongs, as listed at the bottom of the panel. **d**, Geographic distribution of nodes (teams) over the US, color-coded by the structural groups as in panel **c**. The fact that more than one conference is grouped together as groups 3, 4, and 5 can be interpreted in terms of the proximity of the teams' geographic location and its impact on the structure of the network.

approach can be easily adapted to identify network structures defined by link rather than node characteristics[43]. Moreover, it can be applied to networks whose nodes have quantifiable (but not necessarily structural) properties[44], such as age, income and level of education in the case of social networks, which remain elusive in existing network representations. Systematic benchmarking using synthetic networks shows that our method has advan-

tages over existing methods in identifying density-based communities with distinct internal structures (red vs. blue curve in Fig. 3). Naturally, existing methods such as the one proposed in Ref. 14 may still be more effective in resolving specific networks not represented in our benchmarks. In finding general structural groups beyond density-based communities, the visual analytics method outperforms the direct application of standard clustering algorithms in the full node property space (Fig. 1; Supplementary Fig. S1; red vs. green/purple curve in Fig. 3). This suggests that our approach also has potential to be an alternative for solving general high-dimensional clustering problems. The replacement of the human component in the visual analytics method with a simple heuristics based on K-means yields a fully objective unsupervised algorithm, which performs much better than various extensions of K-means directly applied to the full node property space (orange vs. green/purple curves in Fig. 3). This highlights the critical role played by the integrative analysis of clustering outputs from multiple projections. Although the visual analytics method converted to an unsupervised algorithm performs better than standard unsupervised approaches, the original formulation with the human component is still more effective (red vs. orange curve in Fig. 3). By combining the pattern recognition ability of humans with the processing capability of computers, our visual analytics method can resolve the internal organization of complex networks better than either of them alone.

Methods

Biased random projections. To enhance the probability of resolving group separation, we first choose each node property j with probability r_j (while requiring a minimum of four properties) and generate a random projection using those selected properties. The probability r_j is designed to reflect the relative importance of property j in separating the groups. We set $r_j := [\bar{v}_j/\max_j(\bar{v}_j)]^\alpha$, where $\bar{v}_j := \sum_k w_k v_{k,j}^2$, and $v_{k,j}$ denotes the jth component of the normalized basis vector for the kth (out of N) one-dimensional projections generated randomly and uniformly. The weights w_k are given by $w_k := \max_i(z_{k,i+1} - z_{k,i}) \cdot (i/n) \cdot (1 - i/n)$, where $z_{k,1} \le z_{k,2} \le \cdots \le z_{k,n}$ denote the ordered points in the kth projection for all n nodes in the network. The parameter α can be used to adjust the bias strength and was taken to be 2 in all computations.

Controlling for group-size effect in Q_g. Since smaller groups naturally tend to have smaller within-group variations, the ratio of the averages in Eq. (1) increases with the number of groups K, even when the groups are not necessarily better separated. To correct for this bias, we define Q_g by normalizing the ratio by its expected value q_g for randomized groupings with the individual group sizes kept fixed. We estimated q_g by averaging over 100 realizations.

Two-group benchmark networks. For the benchmarking results shown in Fig. 3, we used networks having two groups, constructed as follows. In the larger group (150 nodes), nodes are connected randomly, with the degree of each node fixed to a random integer chosen uniformly between 10 and 70. In the smaller group (50 nodes), node pairs are connected randomly with fixed probability $p_{\rm in}$. Across the two groups, node pairs are connected with probability $p_{\rm out}$. For a given $p_{\rm out}$, we choose $p_{\rm in}$ to match the average degree in the smaller group with the average internal degree in the larger group. The probability $p_{\rm out}$ is varied between 0 (two completely isolated groups) and $40/150 \approx 0.27$ (no internal links in the smaller group), with $p_{\rm out} = 20/150 \approx 0.13$ corresponding to the point at which the average internal and external degrees in the smaller group are equal.

Benchmarking procedure. We used the two-group network described in the subsection above to compare performance of various methods for identifying the groups. For our visual analytics method, we used the node properties listed in Table I and generated 30 biased random projections. The threshold level for the resulting dendrogram was selected so as to produce two groups. In a few cases where a two-group threshold does not exist, we selected the threshold that results in the smallest possible number of groups above two. For the mixture model method[14], the number of groups was set to $K = 2$. For K-means[45], the algorithm was applied directly to the node property space with $K = 2$. For completeness, we also examined the performance of K-means using all possible combinations of choices for (i) kernel[46] (linear, polynomial, Gaussian, or sigmoid); (ii) dimensionality reduction (projecting the data points in the node property space onto the 2, 5, 10, 15, or 20 leading principal components, or no reduction); and (iii) normalization (scaling each node property to have zero mean and unit variance, normalizing each property to the unit interval $[0, 1]$, or no normalization). Scaling for zero mean and unit variance is equivalent to weighing each node property equally when measuring distances in the node property space, while normalizing to the unit interval ensures that all the node properties are distributed in the same range. For the unsupervised variant of our visual analytics method, the human user was replaced by the (linear) K-means algorithm with $K = 1, 2, \ldots, 10$ to analyze each two-dimensional projection, with an optimal choice of K determined by the gap statistic[47], which is defined based on a characteristic signature in the K-dependence of the within-group variation. The performance of each method was measured by the adjusted Rand index R between the computed and the true groupings (see the subsection below for definition).

Rand index. This index measures the similarity between two ways of grouping a given set of discrete objects, possibly into different numbers of groups. For a given pair of groupings of network nodes, the adjusted Rand index R is defined as the normalized fraction of

node pairs that are either classified in the same group in both groupings or classified in different groups in both groupings[38]. The normalization implies that $R = 1$ for identical groupings and $R \approx 0$ for a pair of random groupings.

Ranking node properties. For a given node grouping, we seek a two-dimensional projection that maximizes $\langle n_k ||\mathbf{c}_k - \mathbf{c}||^2\rangle_k / \langle ||\mathbf{x}^{(i)} - \mathbf{c}_{k_i}||^2\rangle_i$, a group separation measure similar to that in Eq. (1) but computed for the projected points after the groups have been identified. Here n_k denotes the number of nodes in group k, and $\mathbf{c}$ denotes the center of all the data points. Such a projection plane can be efficiently found by a spectral method[48] based on the QR decomposition. The node properties are then ranked in the order of increasing angle between their coordinate axes and the projection plane.

Software

A version of the visual analytics software that implements our method for all the networks discussed in this article is available at http://purl.oclc.org/net/find_structural_groups

Acknowledgments

This work was supported by NSF DMS/FODAVA Grant No. 0808860.

Author Contributions

T.N. and A.E.M. designed the research, performed the research, and wrote the manuscript.

* Correspondence and requests for materials should be addressed to T.N. (tnishika@clarkson.edu).

[1] Strogatz, S. H. Exploring complex networks. *Nature* **410**, 268–276 (2001).

[2] Girvan, M. & Newman, M. E. J. Community structure in social and biological networks. *Proc. Natl. Acad. Sci. USA* **99**, 7821–7826 (2002).

[3] Guimera, R., Sales-Pardo, M. & Amaral, L. Modularity from fluctuations in random graphs and complex networks. *Phys. Rev. E* **70**, 025101 (2004).

[4] Radicchi, F., Castellano, C., Cecconi, F., Loreto, V. & Parisi, D. Defining and identifying communities in networks. *Proc. Natl. Acad. Sci. USA* **101**, 2658–2663 (2004).

[5] Palla, G., Derenyi, I., Farkas, I. & Vicsek, T. Uncovering the overlapping community structure of complex networks in nature and society. *Nature* **435**, 814–818 (2005).

[6] Reichardt, J. & Bornholdt, S. Statistical mechanics of community detection. *Phys. Rev. E* **74**, 016110 (2006).

[7] Danon, L., Diaz-Guilera, A. & Arenas, A. The effect of size heterogeneity on community identification in complex networks. *J. Stat. Mech.-Theory E.* **2006**, P11010 (2006).

[8] Fortunato, S. & Barthelemy, M. Resolution limit in community detection. *Proc. Natl. Acad. Sci. USA* **104**, 36–41 (2007).

[9] Chauhan, S., Girvan, M. & Ott, E. Spectral properties of networks with community structure. *Phys. Rev. E* **80**, 056114 (2009).

[10] Porter, M. A., Onnela, J. P. & Mucha, P. J. Communities in networks. *Notices Amer. Math. Soc.* **56**, 1082–1097 (2009).

[11] Chen, P. & Redner, S. Community structure of the physical review citation network. *J. Informetr.* **4**, 278–290 (2010).

[12] Fortunato, S. Community detection in graphs. *Phys. Rep.* **486**, 75–174 (2010).

[13] Mucha, P. J., Richardson, T., Macon, K., Porter, M. A. & Onnela, J. P. Community structure in time-dependent, multiscale, and multiplex networks. *Science* **328**, 876–878 (2010).

[14] Newman, M. E. J. & Leicht, E. A. Mixture models and exploratory analysis in networks. *Proc. Natl. Acad. Sci. USA* **104**, 9564–9569 (2007).

[15] Ravasz, E., Somera, A., Mongru, D., Oltvai, Z. & Barabási, A. Hierarchical organization of modularity in metabolic networks. *Science* **297**, 1551–1555 (2002).

[16] Sreenivasan, S., Cohen, R., López, E., Toroczkai, Z. & Stanley, H. E. Structural bottlenecks for communication in networks. *Phys. Rev. E* **75**, 036105 (2007).

[17] Costa, L. da F., Villas Boas, P. R., Silva, F. N. & Rodrigues, F. A. A pattern recognition approach to complex networks. *J. Stat. Mech.* **2010**, P11015 (2010).

[18] Bagrow, J. & Bollt, E. Local method for detecting communities. *Phys. Rev. E* **72**, 046108 (2005).

[19] Clauset, A. Finding local community structure in networks. *Phys. Rev. E* **72**, 026132 (2005).

[20] Raghavan, U. N., Albert, R. & Kumara, S. Near linear time algorithm to detect community structures in large-scale networks. *Phys. Rev. E* **76**, 036106 (2007).

[21] Rosvall, M. & Bergstrom, C. T. An information-theoretic framework for resolving community structure in complex networks. *Proc. Natl. Acad. Sci. USA* **104**, 7327–7331 (2007).

[22] Hofman, J. M. & Wiggins, C. H. Bayesian approach to network modularity. *Phys. Rev. Lett.* **100**, 258701 (2008).

[23] Kovacs, I., Palotai, R., Szalay, M. & Csermely, P. Community landscapes: An integrative approach to determine overlapping network module hierarchy, identify key nodes and predict network dynamics. *PLoS ONE* **5**, e12528 (2010).

[24] Estrada, E. Community detection based on network com-

municability. *Chaos* **21**, 016103 (2011).

[25] Lancichinetti, A., Radicchi, F., Ramasco, J. J. & Fortunato, S. Finding statistically significant communities in networks. *PLoS ONE* **6**, e18961 (2011).

[26] Psorakis, I., Roberts, S., Ebden, M. & Sheldon, B. Overlapping community detection using bayesian non-negative matrix factorization. *Phys. Rev. E* **83**, 066114 (2011).

[27] Wen, H., Leicht, E. A. & D'Souza, R. M. Improving community detection in networks by targeted node removal. *Phys. Rev. E* **83**, 016114 (2011).

[28] Thomas, J. J. & Cook, K. A. A visual analytics agenda. *IEEE Comput. Graph.* **26**, 10–13 (2006).

[29] Keim, D., Mansmann, F., Schneidewind, J., Thomas, J., & Ziegler, H. Visual Analytics: Scope and Challenges, in *Visual Data Mining*, eds. Simoff, S. J., Bhlen, M. H. & Mazeika, A., Vol. 4404 of *Lec. Notes Comput. Sc.*, 76-90 (Springer, Berlin/Heidelberg, 2008).

[30] von Ahn, L. & Dabbish, L. Labeling images with a computer game, in *Proceedings of the SIGCHI conference on human factors in computing systems*, 319–326 (ACM, 2004).

[31] von Ahn, L., Maurer, B., McMillen, C., Abraham, D. & Blum, M. reCAPTCHA: Human-based character recognition via web security measures. *Science* **321**, 1465–1468 (2008).

[32] Jain, A., Murty, M. & Flynn, P. Data clustering: A review. *ACM Comput. Surv.* **31**, 264–323 (1999).

[33] Donetti, L. & Muñoz, M. A. Detecting network communities: a new systematic and efficient algorithm. *J. Stat. Mech.* **2004**, P10012 (2004).

[34] Seary, A. J. & Richards, W. D. Partitioning networks by eigenvectors, in *Proceedings of the International Conference on Social Networks*, Vol. 1, 47–58 (1995).

[35] Pothen, A., Simon, H. D. & Liou, K. P. Partitioning sparse matrices with eigenvectors of graphs. *SIAM J. Matrix Anal. Appl.* **11**, 430–452 (1990).

[36] Newman, M. E. J. Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E* **74**, 036104 (2006).

[37] Chung, F. *Spectral Graph Theory* (American Mathematical Society, Province, 1997).

[38] Hubert, L. & Arabie, P. Comparing partitions. *J. Classif.* **2**, 193–218 (1985).

[39] Estrada, E. & Rodríguez-Velázquez, J. A. Subgraph centrality in complex networks. *Phys. Rev. E* **71**, 056103 (2005).

[40] Zachary, W. W. An information flow model for conflict and fission in small groups. *J. Anthro. Res.* **33**, 452–473 (1977).

[41] Thiemann, C., Theis, F., Grady, D., Brune, R. & Brockmann, D. The structure of borders in a small world. *PLoS ONE* **5**, e15422 (2010).

[42] Ratti, C., Sobolevsky, S., Calabrese, F., Andris, C., Reades, J., Martino, M., Claxton, R. & Strogatz, S. H. Redrawing the map of Great Britain from a network of human interactions. *PLoS ONE* **5**, e14248 (2010).

[43] Ahn, Y. Y., Bagrow, J. P. & Lehmann, S. Link communities reveal multiscale complexity in networks. *Nature* **466**, 761–764 (2010).

[44] Bianconi, G., Pin, P. & Marsili, M. Assessing the relevance of node features for network structure. *Proc. Natl. Acad. Sci. USA* **106**, 11433–11438 (2009).

[45] Lloyd, S. P. Least squares quantization in PCM. *IEEE T. Inform. Theory* **28**, 129–137 (1982).

[46] Schölkopf, B., Smola, A. & Müller, K. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Comput.* **10**, 1299–1319 (1998).

[47] Tibshirani, R., Walther, G. & Hastie, T. Estimating the number of clusters in a dataset via the gap statistic. *J. Roy. Stat. Soc. B* **32**, 411–423 (2001).

[48] Park, H., Drake, B., Lee, S. & Park, C. Fast linear discriminant analysis using QR decomposition and regularization. *Georgia Institute of Technology, GA, Tech. Rep. GT-CSE-07-21* (2007).

[49] Krebs, V. http://www.orgnet.com/

[50] Goh, K. I., Cusick, M. E., Valle, D., Childs, B., Vidal, M. & Barabási, A. L. The human disease network. *Proc. Natl. Acad. Sci. USA* **104**, 8685–8690 (2007).

[51] Gürsoy, A. & Atun, M. Neighbourhood preserving load balancing: A self-organizing approach, in *Euro-Par 2000 Parallel Processing*, eds. Bode, A., Ludwig, T., Karl, W. & Wismüller, R., Vol. 1900 of *Lec. Notes Comput. Sc.*, 234–241 (Springer, Berlin/Heidelberg, 2000).

Supplementary Information

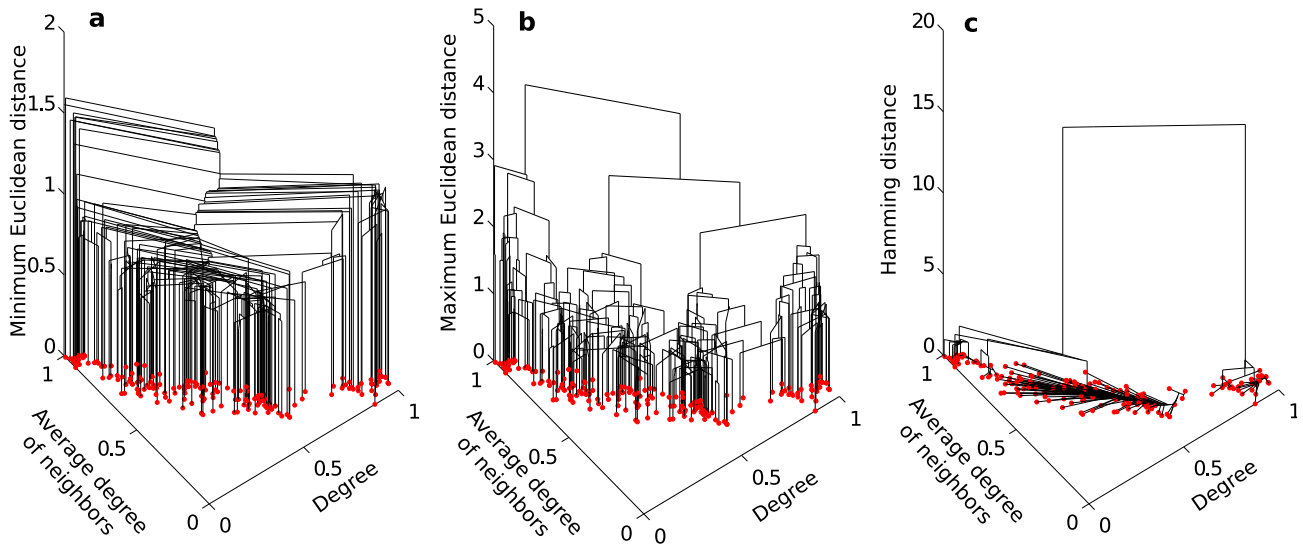

Supplementary Figure S1: **Comparing hierarchical clustering methods. a**, **b**, Dendrograms obtained by applying the single-linkage (**a**) and complete-linkage (**b**) clustering algorithms[32] to the full 28-dimensional node property space. We use an example network with two structural groups defined by their degree distributions, which we construct by randomly connecting nodes with prescribed degrees. A group of 150 low-degree nodes (degree between 10 and 50) and a group of 30 high-degree nodes (degree between 70 and 90) form a bimodal degree distribution, bridged by 5 nodes with degree between 50 and 70. All node degrees are chosen uniformly at random from the corresponding intervals. The nodes (red dots) are plotted in the two-dimensional projection using the degree and the average degree of neighbors, both normalized to the unit interval. For these two dendrograms, the distance between two nodes is measured by the Euclidean distance in the full node property space. In the single-linkage dendrogram (**a**) the height at which two groups join is the minimum distance between all pairs of nodes from the two groups, while in the complete-linkage dendrogram (**b**) it is the maximum distance between node pairs. **c**, Dendrogram obtained by our visual analytics method. In applying single-linkage clustering to obtain the dendrogram, the measure used for distance between nodes i and j is the Hamming distance between the user input vectors $\mathbf{a}^{(i)}$ and $\mathbf{a}^{(j)}$ (which can be defined in this case, but not in the other two cases). Two or more groups joined in this dendrogram at height d implies that any pair of nodes taken from two of these groups has been separated visually by the user in at least d different projections. We see that thresholding at a fixed level in this dendrogram would clearly produce the correct splitting of nodes into the two structural groups, while the dendrograms in the other two panels do not capture the group structure. This comes from the difficulty in recognizing groups with wide distribution of points, also suffered by the K-means algorithm (see Fig. 1**d**).

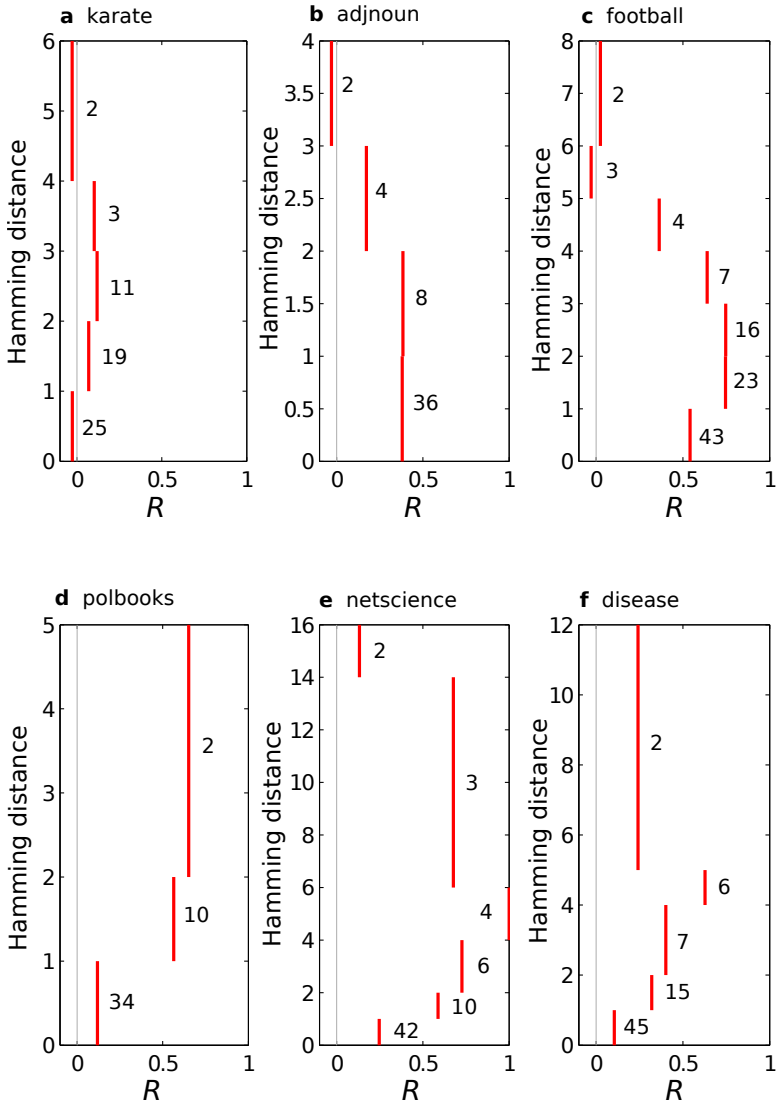

Supplementary Figure S2: **Comparing discovered groups with communities found by a traditional method.** For each network in Table II, we compare the grouping obtained at a fixed hierarchical level in our method with the division into the same number of groups that is found by the method of Ref. 2, which is based on link betweenness centrality. The similarity of the two groupings is measured by the adjusted Rand index R (defined in Methods). Each panel shows R as a function of the hierarchical level measured by the Hamming distance in our method, with the corresponding number of groups indicated in the plots. The R values significantly lower than unity are observed at most hierarchical levels for each network, confirming that the group structures we discovered are indeed different from the traditional community structures.

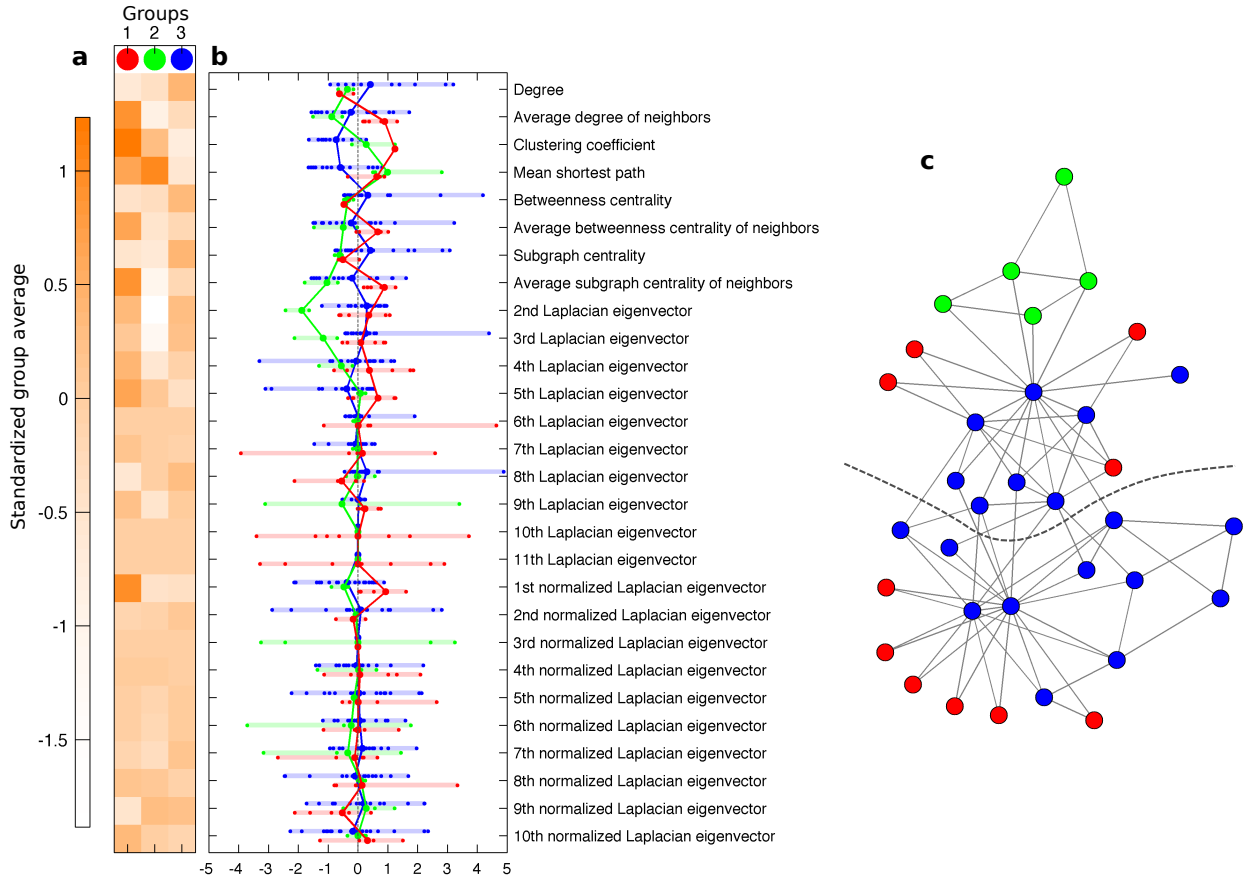

Supplementary Figure S3: **Characterizing three structural groups discovered in the karate network.** **a**, Average node properties of the three groups. Rows correspond to node properties and columns to groups (red, green, and blue disks at the top). Using the orange color-scale on the left, each cell shows the average node property of the group, relative to the network average and in units of the network standard deviation. **b**, Node property distribution within each group. The three groups in this plot are color-coded as in the disks at the top of panel **a**. Small dots indicate the individual values for each node in the network, larger dots connected by lines indicate the group averages, and bars indicate the range of values for each group. All values are measured relative to the network average and in units of the network standard deviation. **c**, Layout of the network with groups color-coded as in the other panels. The dashed curve indicates the eventual split of the karate club into two different clubs, as documented in Ref. 40. It is clear that the group structure discovered by our visual analytics approach is distinct from the club split. Group 1 (red) is characterized by low degree, clustering coefficient of one, and large values of the components of the first normalized Laplacian eigenvector, as well as by being connected to high-degree nodes with high betweenness and subgraph centrality (the first column of panel **a** and red plots in panel **b**). Group 2 (green) is characterized by very low values of the components of the second Laplacian eigenvector, which is indicative of a traditional community structure. This is reflected in panel **c**, where green nodes form a cluster at the top. High values of mean shortest path distance to other nodes for nodes in Groups 1 and 2 implies that these nodes sit in peripheral locations within the network, which is confirmed in the network layout in panel **c**. Group 3 (blue) forms the core of the network and includes all the high-degree, high-centrality nodes, which is reflected in the high group average in these quantities. This example network illustrates that our method identifies groups of nodes with common structural properties that do not even need to be connected within each group.

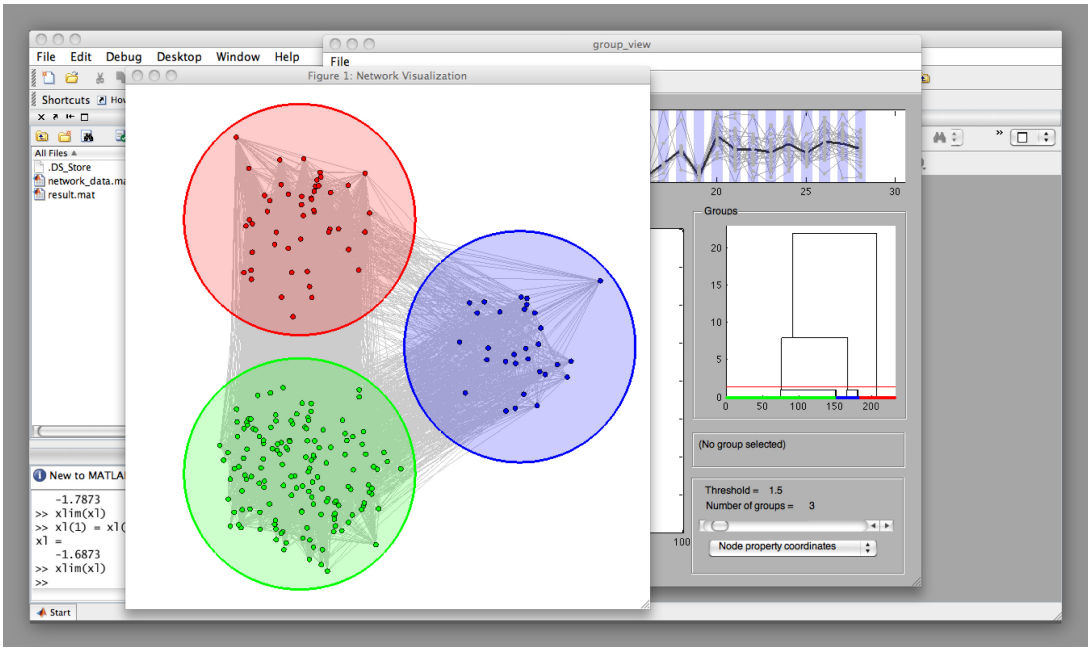

Click on the image to see the video on the *Scientific Reports* website (MOV Format, Length 4:47).

Supplementary Video S1: **Visual analytics software for discovering structural groups in networks.** The network shown in Fig. 1 is used to demonstrate how an implementation of the method introduced in this article discovers the three groups of nodes characterized by the degree and the average degree of neighbors. The default set of node properties is the same as in Table I, and we use 30 random projections with the bias described in Methods to enhance the probability that the user sees group separation. Only a selection of these projections are actually shown in the movie in order to keep the presentation short.

Supplementary Table SI: **Description of networks analyzed.**

Dataset		Description	Reference
karate		Social network of a university-based karate club	40
	Node:	A club member	
	Link:	Interaction between the two members in at least one context outside the club activity	
polbooks		Network of books on American politics bought from Amazon.com	49
	Node:	A book	
	Link:	Frequent purchase of the two books together by the same buyer	
adjnoun		Network of nouns and adjectives appearing in a novel (*David Copperfield* by Charles Dickens)	36
	Node:	A noun or adjective	
	Link:	Appearance of the two words adjacent to each other in the book	
football		Network of collegiate American football teams in the US	2
	Node:	A football team	
	Link:	The fact that one or more games were played between the two teams in the 2000 regular season	
netscience		Largest connected component of the network of scientists who have published papers on network science	36
	Node:	A scientist	
	Link:	Coauthorship between the two scientists	
disease		Largest connected component of the network of known human genetic disorders	50
	Node:	A genetic disorder	
	Link:	Existence of a common gene whose mutation is associated with both disorders	